

10/587618

EQ 619695093 VS

AP20 Rec'd PCT/PTO 27 JUL 2006

HIERARCHICAL RESOURCE MANAGEMENT FOR A COMPUTING UTILITY

CROSS REFERENCES

The present application is cross-referenced to application entitled "Componentized Automatic Provisioning and Management of Computing Environments for Computing Utilities", having docket number YOR920030588PCT1, and application entitled "Apparatus for Arbitration in a Computing Utility System", having docket number YOR920030589US1, even dated herewith, and which are included herein by reference in entirety for all purposes.

FIELD OF THE INVENTION

The present invention is directed to provisioning and management of resources in a computing utility, which uses the resources to provide computing services to customers. It is more particularly directed to hierarchical resource management for the computing utility.

BACKGROUND

A hosted center provides computing services to multiple customers. Each customer is allocated a subset of the service provider's infrastructure resources, such as servers, to meet its needs. Customer needs change over time, and in particular, peak demands for resources can exceed average demands by orders of magnitude. A simple static allocation of resources to satisfy peak demand leaves resources underutilized. Reconfiguring infrastructure resources dynamically, in response to customer needs, demands prompt attention from administrative personnel and could require moving hardware, increasing both operational costs and the risk of failing to provide

adequate service. The problem for the hosted center is how to respond quickly to changes in customer needs such that infrastructure resources and staff are employed in an efficient and cost effective way. *Computing utilities* attempt to address this problem by automating the creation and management of multiple computing services on shared dynamically allocatable infrastructures.

Previous work in computing utilities varies in the types of services that are offered, the resources that are used, and the extent to which operation is automated. Operations subject to automation are wide ranging, and include creating services, deploying a service for a customer, modifying the set of resources used to provide the service, and incorporating new resource instances and types into the hosted center and its services.

A previous system automated the provisioning of front end servers in web sites based on metrics such as server load and response time. It included a component that discovered server and network topology automatically. Another system provided a variety of multi-tier web sites in which pre-configured servers could be allocated to different tiers automatically based on metrics such as server load. Yet another system also allocated server resources in response to server load, but modeled both the value of allocating resources to each customer and the cost of employing those resources, with an emphasis on energy cost. More recent work includes allocation of other resource types such as memory and storage, and allocation of servers for general use.

At the application layer some systems have a framework for deployment and management of distributed applications. An application is described as a collection of related, reusable components, which may represent resources or subsystems. The description includes dependency information to ensure that, for example, components are started in the correct sequence. Once deployed, applications may be monitored, and actions may be specified in case of component or resource failures, such as automatic fail over or restart. Such a system is not used for low level resource configuration tasks such as installing operating systems on servers, but for higher level application specific configuration.

A growing number of industrial products aim to provide multi-tier applications over a physical

1 infrastructure consisting of a variety of resources, such as those from Hewlett Packard,
2 ThinkDynamics, Sun Microsystems, and Jareva. They vary in many respects, such as the types of
3 resources provided (e.g., servers and storage); specific operating systems and middleware
4 supported; assumptions and characteristics of the network infrastructure (e.g., whether or not
5 network isolation is provided via VLAN); level of monitoring support (e.g., resource usage,
6 failure detection, SLA, threshold based alerting); support for resource discovery; support for
7 modifying service resources once deployed; whether modifications can occur automatically (e.g.,
8 triggered by SLAs); and the extent to which the products can or must be customized to fit
9 preexisting hosted center infrastructures.

10 **SUMMARY OF INVENTION**

11 This invention provides methods, apparatus, systems and architectures for a hierarchical
12 provisioning and management of a computing infrastructure, used to provides service oriented,
13 or utility computing. It is useful for hosted environments. It is often provided to customers of a
14 service provider who owns an infrastructure. It allows the dynamic provisioning and assignment
15 of resources to computing environments.

16 An aspect of the present invention is to provide an element of a computing utility that operates
17 in an environment that possesses the following characteristics. Together, these characteristics
18 generalize the environments studied in previous work.

19 Another aspect of the present invention is to provide for automatically provisioning and
20 managing resources for a computing utility. A computing utility could be used by a corporation,
21 a service provider, or an individual.. This invention allows the sharing of resource within sets of
22 customers and allows for restricting resources to specific customers. It allows customers to offer
23 one or more services composed of managed resources to their clients. Managed resources may be
24 assigned to customers or domains on demand. It also allows resources to be hierarchically
25 managed.

1 In an example embodiment of a method according to the present invention, the method
2 comprises the step of providing hierarchical management of at least one domain for an entity.
3 The step of providing hierarchical management including: obtaining a hierarchical representation
4 of the at least one domain, the representation including: a list of computing environments to be
5 managed, at least one policy controlling acquisition of at least one resource from resource
6 libraries for the at least one domain, and any sub-domains within the at least one domain; and
7 instantiating the representation.

8 In an example embodiment of an architecture according to the present invention, the architecture
9 for a computing utility comprises an apparatus to provide at least one service for a plurality of
10 clients. The apparatus comprises: a Base Resource Distribution Service to allocate resources to
11 the at least one service; the Base Resource Distribution Service having at least one collector; at
12 least one Provisioned and Managed Resource Service coupled to the Base Resource Distribution
13 Service to provision and manage the resources for the at least one service; and at least one Base
14 Resource Library Service coupled to the Base Resource Distribution Service to provide
15 reservation and allocation of resources.

16 **BRIEF DESCRIPTION OF THE DRAWINGS**

17 These, and further, aspects, advantages, and features of the invention will be more apparent from
18 the following detailed description of a preferred embodiment and the appended drawings
19 wherein:

20 Fig. 1 shows the components of a Hosted Environment;

21 Fig. 1A shows an example structure of a company;

22 Fig. 2 shows the components in a computing utility management system in accordance with
23 the present invention;

24 Fig. 3 shows the operations of a library component in a computing utility management system

1 in accordance with the present invention;

2 Fig. 4 shows a the hierarchy of collectors within a computing utility in accordance with the
3 present invention;

4 Fig. 4 A shows applications allocated to a unit of a company in accordance with the present
5 invention;

6 Fig. 4 B shows how the applications of Fig. 4A can be structured for provisioning in
7 accordance with the present invention;

8 Fig. 5 shows a computing utility with multiple root collectors in accordance with the present
9 invention;

10 Fig. 6 shows the association of resource pools to collections in accordance with the present
11 invention;

12 Fig. 7 show how acquisition policy is checked within a computing utility in accordance with
13 the present invention;

14 Fig. 8 show how to locate available resources within a computing utility;

15 Fig. 9 shows the interaction between a Provisioned and Managed Resource Service (PMRS)
16 and a Base Distribution Resource Service (BRDS) in a computing utility in accordance with
17 the present invention;

18 Fig. 10 show how a collector functions during the operation of a computing utility in
19 accordance with the present invention;

20 Fig. 11 Illustrates how a new service is added to a computing utility in accordance with the
21 present invention;

22 Fig. 12 Shows applications assignment within a company in accordance with the present
23 invention; and

24 Fig. 13 shows the collector hierarchy associated with the assignment of Fig. 12, in accordance
25 with the present invention.

26 **DETAILED DESCRIPTION OF THE INVENTION**

27 This invention provides methods, apparatus, systems and architectures for hierarchical

1 provisioning and management of a computing infrastructure which is used to provides service
2 oriented, or utility computing, to the customers of the service provider who owns the
3 infrastructure. In this way hosted environments can acquire resources from other hosted
4 environments, etc. The invention provides an architecture for hierarchical management of
5 computing infrastructures which included arbitration, provisioning and management. It allows
6 the dynamic provisioning and assignment of resources to computing environments. Customers
7 can have multiple computing environments within their domain. The computing utility shares its
8 resources across multiple customer domains and arbitrates on the use of resources between and
9 within domains. The invention enables resources to be dedicated to a specific customer domain
10 or to a specific computing environment. Customers can specific acquisition and distribution
11 policy which controls their use of resources from the utility within their domains.

12 The present invention is an element of a computing utility that generally, but not necessarily,
13 operates in an environment that possesses one or more of the following characteristics. Together,
14 these characteristics generalize environments studied in previous work.

15 First, resources may be allocated to customers in combinations which are heterogeneous,
16 may be interdependent, and vary over time.

17 Second, the services provided to each customer may be different. For example, one
18 customer may be provided resources for a web site, and another for a scientific computing
19 cluster. Resource types, quantities, dependencies, and allocation patterns will thus vary
20 between customers.

21 Third, the level of service provided to each customer may be different. This means that
22 evaluating the quality of a resource allocation to a customer considers both the type and
23 level of service.

24 Fourth, the resource infrastructure varies between service providers. Further, for a given
25 service provider, the infrastructure varies over time. These variations can be a result of

1 upgrades or additions to the physical infrastructure, engaging the services of another
2 provider, or inclusion of additional resources when idle or at certain times of day, or
3 removal of resources which are no longer needed.

4 Fifth, resources may be allocated in advance, or reserved. A reservation may be accepted
5 even if there are not sufficient resources currently available in the infrastructure to fulfill
6 it, based on the expectation that additional resources will be acquired in time for
7 allocation.

8 Sixth, customers may want to share resources within an organization by subdividing their
9 allocations and having the service provider manage the subdivisions, such as to
10 departments.

11 Seventh, customers may want to supply resources which they currently own to be
12 managed by the service provider (i.e., the customer outsources the management of its
13 resources to the service provider), which can only be used for their purposes. This
14 requirement imposes constraints on where resources may be allocated.

15 Eighth, customers may wish to specify policies which govern the management and
16 operation of their resources. These policies are to be implemented by the computing
17 utility.

18 Finally, different service providers may have different objectives in mind when allocating
19 resources to customers, such as maximizing profit, availability, or performance.

20 This invention is also useful as a part of an apparatus for automatically provisioning and
21 managing resources for a computing utility. A computing utility could be used by a corporation, a
22 service provider, or an individual.. This invention allows the sharing of resource within sets of
23 customers and allows for restricting resources to specific customers. It allows customers to offer
24 one or more services composed of managed resources to their clients. Managed resources may be

1 assigned to customers or domains on demand. It also allows resources to be hierarchically
2 managed.

3 Figure 1 illustrates a Hosted Environment where the present invention runs. This environment is
4 made up of processors 101, storage 103, firewalls 105, and software 107. The software 107 can
5 be operating systems, middleware or applications. In Figure 1 the software 107 is represented as
6 Business Processes, Business Applications, or Services. All of the available software has been
7 pre-configured into these elements. In fact it is the purpose of this invention to enable the
8 dynamic reconfiguration of the hardware and software components into new or existing services
9 as needed. This type of environment can exist within a large company or it can be offered by an
10 ISP or ASP as a basis for service. Returning to the detailed description, the hardware resources
11 are connected by a network as indicated by the grid of lines 109 interconnecting all of these
12 resources. This network may be configured into one or more tiers, where each tier is separated by
13 a router 105 or firewall 105. The configuration of the tiers within the network may be static or
14 dynamic. When a network is statically configured it is not possible without human (or
15 mechanical) intervention to move a resource between tiers in the network. With a dynamic
16 configuration, resources can be moved between different tiers under the control of an
17 infrastructure, such as the one being described in this invention, or an operator working at a
18 console. Within each tier Resources can be arbitrarily grouped into units. This can be done by
19 using mechanisms such as Virtual LANs. Software resources are assigned to physical resources
20 by the controlling infrastructure.

21 In this environment a subset of the resources are assigned to the management infrastructure 111,
22 113, & 115. In Figure 1 these resources are indicated by a dotted line surrounding them. The
23 resources assigned to the management infrastructure run the Resource Management Software
24 described in this invention. This software manages the rest of the resources. The resources used
25 for the management infrastructure are not assigned to clients 117 of the hosted environment. The
26 managed resources are assigned to clients 117 as necessary. It is expected that clients 117 will
27 primarily receive service by connecting to the hosted environment through the Internet 119.
28 However, clients 117 can receive service if they are connected by any means to the hosted

environment. For example, they could have a direct connection to a managed resource, they could be connected to the same network as the managed resources, they could have a VPN connection to the managed resource, or they could have a VPN connection to the network containing the managed resources. Clients 117 have no view to the infrastructure; they only see the applications or services that they are receiving from the hosted environment.

The management infrastructure assigns resources to each supported service as needed. The need is determined by the demand on the service in accordance with SLAs, contract, or other agreement that authorized the service to operate in the hosted environment. Individuals or organizations that wish to offer services will have agreements with the provider of the hosted environment. Returning to Figure 1, the processors 101 may have direct attached storage that would be managed as part of managing the processor. Storage 103 in this Figure refers to any form of storage server that can be network attached. The processors can also represent some complex function like a database server or a video server. The degree to which a resources is divided into smaller functions for management is dependent on the owner of the infrastructure. Even if a complex resource such as a database server can be broken down into smaller components, the owner of the resource will determine how it is represented within the management infrastructure. This feature allows the invention to use resources that are not local to the hosted environment. The invention assumes that the resources are securely interconnected with adequate performance. However, the network which connects the resource represents a logical connection. Any infrastructure, such as a VPN, which allows an adequate degree of security, performance, and control over a resource, can be used to connect resources into the hosted environment. Consequently the hosted environment as depicted in Figure 1 could itself be distributed world wide. The Internet in this picture represents a logical placement the Internet can be used to dynamically add resources to the hosted environment.

A **domain** is an organizational unit within a customer for which services are provisioned. An organization may have multiple sub domains, division, units, or departments within it which also utilize the service provider. A sub domain is also a domain and the domains of an organization are form a tree. An organization can have multiple trees of domains which are not connected. For

1 the purposes of this invention, we consider those to be multiple organizations. The root domain
2 is the highest domain or the base of the domain tree for an organization. Figure 1A shows the
3 domains within a customer where services can be provided. In Figure 1A the main company, or
4 root domain, is Smith Fastener Company 120, all other parts of this company are subdivisions,
5 or sub-domains, of Smith Fastener. Smith Fastener has four main divisions, Research 121,
6 Finance 122, Hardware 123, and Marketing 124. The hardware and marketing division have
7 additional units, or sub domains, within each of them. Hardware is divided into Bolts 127, screws
8 128, and hinges 129. Marketing has two regions Region 1, 125 and Region 2, 126. Region 2 is
9 further divided into North 130 and South 131. It is an object of the present invention to allow
10 service providers to provide computing services to the various domains, units, and/or division of
11 any company, such as the Smith Fastener company. Any domain or sub-domain may then
12 subscribe to the computing services provided. This invention assumes that for provisioning
13 purposes the company can be structured into a domain tree. Further any arbitrary tree can be
14 supported. As will be described later, the domain tree controls the sharing of resources for the
15 purposes of provisioning between the various applications that are used by a company. This
16 invention does not require that the resource sharing tree exactly map to the corporate structure.
17 Any arbitrary resource sharing tree can be used. For illustrative purposes we will use one that
18 maps to the structure of an organization, Smith Fastener Company 120.

19 Within an organization, each domain, division, or unit has its own employees. This invention
20 does not restrict who (which employees of the corporation) within a corporation can use the
21 computing services. Eligibility to use a computing environment is controlled by the
22 administrative policies within the corporation. User Ids, passwords, and authorization are
23 controlled by the corporation. For example, all of the application could use a common
24 authentication and authorization mechanisms, each application could have its own, or any
25 combination in between. The provisioning of access control is part of the PMRS. The use, (user
26 who are allowed to access or use the applications) of computing services is independent of the
27 structure used to provision resources to the applications.

28 In this environment there are multiple ways to divide the roles and responsibilities for utilizing

1 this invention. In one such division there is an infrastructure provider who owns the hosted
2 environment described in Figure 1; one or more Service Providers offering some service on
3 using the hosted environment; and one or more clients of the Service Provider. Note that the
4 service provider is also a customer of the infrastructure provider. In another such division a
5 company (or entity) may contract with an infrastructure provider for services, outsourcing the
6 management of its infrastructure. In this case the company is utilizing this invention to manage
7 the services provided within itself. In another such division the infrastructure provider and
8 service provider could be the same. Further there are entities that develop the services that are
9 offered. They could be separate, such as an independent software vendor (ISV), or part of one or
10 more of the previously named entities. An entity who wants to offer a service that is managed by
11 this invention on some infrastructure provides a description of that service to the infrastructure.
12 This description can be abstract or concrete and includes the organization of the service, an
13 indication of the service to be provided and the policies which control the operation of the
14 service.

15 Figure 2 shows a high level view of a computing utility. There are three layers of components.
16 The top layer represents the services which are provisioned, called **computing environments**. A
17 computing environment includes hardware and software resources such as servers, operating
18 systems, applications, and middleware. Each computing environment has a service which is
19 invoked to provision and manage its resources called the Provisioned and Managed Resource
20 Service (PMRS). The PMRS is described in detail in a separate invention entitled
21 "Componentized Automatic Provisioning and Management of Computing Environments for
22 Computing Utilities", having docket number YOR920030588PCT1. The bottom layer
23 represents the resources of the hosted center available to its customers. A **base** resource is a
24 resource that is atomic; it is not broken down into other resources. Each base resource type has a
25 service that provides reservation and allocation of instances of that resource type, called the Base
26 Resource Library Service (BRLS). The definition of base resource is up to the service provider.
27 Therefore, a BRLS can provide a simple resource such as a software license or a complex
28 resource such as an operating system installed and running on a hardware platform. Adding,
29 removing, or modifying a base resource changes the overall capacity of the system. This

invention assumes that both the number and types of the base resources change over time. Base resources are represented by a set of **Base Resource Library Services (BRLS) 201**. Figure 2 shows four types of base resources: DB2 license 203, zSeries LPAR 205, xSeries Server 207, and AIX license 209. The BRLS provides operations such as cataloging, check out (allocation), check in (deallocation), and reservation. The interfaces to a BRLS 203, 205, 207, 209 are described in more detail in Figure 3.

A **composite** resource is built out of one or more other resources (including other composite resources) to perform a specified function. A composite resource may have a set of **dependencies** associated with it. An example of a composite resource is a web site. It may be composed of some number of front-end servers, back end servers, a load balancer, a set of IP addresses for the servers, web server software, database software, and licenses associated with the software. The set of base resources used to implement the function of a composite resource can change over time, though not all composite resources may have this capability. The computing environment is a type of composite resource.

Associated with each computing environment is software that provisions and manages the resources used to provide the service, called the **Provisioned and Managed Resource Services (PMRS) 211**. Every resource type, including composite resources, has a PMRS which encapsulates the knowledge of how to create a resource of that type, and how to manage an instance of the resource. Provisioning refers to the act of allocating a resource to a computing environment and configuring it for use within the service. Provisioning actions may include assembling base resources into composites, configuring network devices, installing operating systems, application software, monitors, and user accounts. Managing the resource may include such actions as requests for adding or removing capacity from the resource instance. Figure 2 shows the PRMSs, DB2 213, Linux 215, and Web Site 217. The PMRS is described in detail in invention entitled "Componentized Automatic Provisioning and Management of Computing Environments for Computing Utilities", having docket number YOR920030588PCT1.

The middle layer of Figure 2 is the **Base Resource Distribution Service (BRDS) 219**. Its job is

to allocate the resources of the service provider to its domains in an effective way, based on the agreements that the service provider has in force for each of its domains. Domains are hierarchically organized, some domains are sub domains of others. The BRDS specifies which sets of resources are available to which domains and automatically redistributes those resources based on business policies when resources are constrained. It interacts with both the computing environments 211 (PMRSs) and resources 201 (BRLSs). A resource may be allocated to a computing environment either immediately or in the future. The promise of a future resource allocation is a **reservation**. This invention assumes that the resources allocated to a computing environment are dedicated to that computing environment for the duration of their allocation.

Base Resource Library Service

Resources reside in **pools**, which may be arranged by type. Resource pools may be public or private. A **public resource pool** is a resource pool from which any domain may allocate a resource. A **private resource pool** is restricted to a subset of customer domains. Private pools are used to dedicate resources to a subset of customer domains. For example, a private pool could hold resources for use in multiple domains on behalf of a single customer. Such a pool might also be used for resources the customer owns but wishes the service provider to manage for it. Every resource has a **home pool** from which it originates, which may be either public or private. The home pool does not change when a resource is allocated or reserved. A group of unallocated resources is called a **free pool**.

Figure 3 shows an example of an operation of a Base Resource Library Service 310 (BRLS). Every base resource in the system is represented by a library, or BRLS 310, which provides information such as cataloging, check out (allocation), check in (deallocation), and advance reservation. A BRLS serves as a home pool for instances of a resource type, and it may be either public or private (represent a public resource pool or a private resource pool). Each BRLS 310 is associated with a collector. Every BRLS 310 has two types of interfaces, Resource Operations 330 which are used by the system to, acquire, reserve, and return resources and Catalogue operations 340 which are used manage the resources available to a BRDS 219. The Resource Operations 330 provided are Reserve, CancelReservation, CheckIn, CheckOut,

1 ExtendReservation, and Query.

2 **Reserve(num-instances, selection-spec, start time, duration)**
3 **-> reservation-ids**

4 This request is issued by BRDS 219 on behalf of a customer computing
5 environment to reserve resources.

6 Inputs:

7 Num-instances is the number of instances desired

8 Selection-spec is a specification of desired attributes supported by the
9 BRLS 310.

10 Start time, If unspecified, any instance may be used to satisfy the request.

11 Start time could be immediate or for an advance reservation

12 Outputs:

13 Reservation-ids[]. The elements are resource reservation tickets, one per
14 resource instance, up to num-instances. These need not map to instances
15 at the time of the reservation. In other words, the library may overbook its
16 resources.

17 **CheckOut(reservation-id) -> resource-handle**

18 This request is issued by BRDS 219 on behalf of a customer computing
19 environment to allocate resources.

20 Inputs:

21 Reservation-id issued by Reserve

22 Outputs:

23 Resource-handle. An identifier of a resource instance. Set to a special
24 value if an instance of the resource cannot be provided.

25 **CheckIn(resource-handle)**

26 This request is issued by BRDS 219 on behalf of a customer computing
27 environment to return resources.

1 Inputs:

2 Resource-handle issued by CheckOut.

3 **ExtendReservation(resource-handle, end time)**

4 This request is issued by BRDS 219 on behalf of a customer computing
5 environment to extend the reservation for a currently held resource.

6 Inputs:

7 Resource-handle issued by CheckOut.

8 End time: the point at which the resource will be returned.

9 Outputs:

10 Accept: Resource reservation can be extended.

11 Reject: Resource is returned.

12 **Query(selection-spec) -> availability-data**

13 This request is issued by BRDS 219 on behalf of a customer computing
14 environment to return resources.

15 Inputs:

16 Selection-spec is a specification of desired attributes supported by the
17 BRLS 310.

18 It is used to restrict the scope of the query.

19 It may contain a resource handle to refer to a specific instance.

20 Outputs:

21 Availability-data[]. The elements are structures that indicate when
22 instances are available for reservation.

23 **CancelReservation(reservation-id)**

24 This request is issued by BRDS 219 on behalf of a customer computing
25 environment to cancel a reservation. After a cancellation is received, the resource is
26 available to be assigned and checked out by another computing environment.

27 Inputs:

1 Reservation-id issued by Reserve

2 *Catalog Operations*

3 Catalog Operations 340 are provided for administrators 320 or other administrative mechanisms
4 to modify the resource pools managed by the BRLS 310. Resource discovery mechanisms could
5 also use these operations to automate this process. The operations provided include Add,
6 Remove, Update, Query, and any combination of these operations.

7 **Add(resource-identifier, instance-data)**

8 This request is issued to add a resource instance to the BRLS 310.

9 Inputs:

10 Resource-identifier is a resource identifier, may or may not be the same as
11 the handle above.

12 Instance-data is the collection of attribute values that can be used to select
13 the resource via a selection-spec

14 **Remove(resource-identifier)**

15 This request is issued to remove a resource instance from the BRLS 310.

16 Inputs:

17 Resource-identifier is the resource identifier

18 **Query(selection-spec) -> resource-identifier[]**

19 This request is used to look up a resource instance registered with the BRLS 310.

20 Inputs:

21 Selection-spec is a specification of desired attributes supported by the
22 BRLS 310.

23 It is used to restrict the scope of the query.

24 It may include a resource identifier.

25 Outputs:

1 Resource-identifier[] a list of resource identifiers that satisfy the
2 selection-spec
3
4 **Query(resource-identifier) -> instance-data**
5 This request is used to look up a resource instance registered with the BRLS 310.
6 Inputs:
7 Resource-identifier is the resource identifier
8 Outputs:
9 Instance-data is the collection of attribute values registered for this
10 instance, if any
11
12 **Update(resource-identifier, update-data) -> instance-data**
13 This request is used to modify a resource instance registered with the BRLS 310.
14 Inputs:
15 Resource-identifier is the resource identifier
16 Update-data new information to be applied to or registered about the
17 resource
18 Outputs:
19 Instance-data is the collection of attribute values registered for this
20 instance, if any

21 *Base Resource Distribution Service*

22 Figure 4 shows an example of components of a Base Resource Distribution service (BRDS) 219.
23 The BRDS 219 determines how to distribute the base resources across the various computing
24 environments. The BRDS 219 includes of two types of components, **Collectors** 420, 422, 424,
25 426, 428 and an **Arbiter** 430. A domain is associated with at least one collector. For every
26 computing environment there is one collector which includes the policies previously mentioned,
27 and the list or resources which are reserved for the computing environment.

Arbiter

The arbiter determines how resources are divided amongst the computing environments. It operates both on current and future allocations (reservations). Resources may be allocated to computing environments either reactively or proactively. The arbiter can be consulted when free resources are not sufficient to fulfill demands from customer computing environments. It can also be used on a periodic basis to optimize resource allocations. If resources are constrained, the arbiter may reclaim resources from computing environments. Details on the operation of the arbiter are in a separate invention entitled "Apparatus for Arbitration in a Computing Utility System", having docket number YOR920030589US1.

Collector

A collector represents the set of resources allocated to one or more computing environments. The root domain of each computing environment has an associated collector. Collectors may be nested, for example, to represent departments within an organization. This structure enables resource sharing between a subset of computing environments, essentially replicating the computing utility model for those domains. A computing utility can provision resources for multiple organizations, a collector that is connected to multiple organizations (the root domain of each organization) is called the **root collector**. Figure 4 shows one root collector 420 with two domains, one of which has two sub domains. The two main domains are Col-A 424 and Col-B 422 representing Company A and Company B. Col-A 424 has two sub domains Col-X 426 and Col-Y 428 representing Departments X & Y. There are three PRMSs shown in Figure 4 PMRS-X 496, PRMS-Y 498, and PRMS-B 495. A collector is associated with each of these computing environments. In addition the BRDS 219 also has a collector 420.

Figure four also shows public and private resource pools. Public BRLSs are associated with a root collector in a BRDS 219. A private BRLS is associated with a collector other than a root collector of a BRDS 219. In Figure 4, BRLSs 486, 488, and 482 are public because they are associated with the root collector 420 and BRLS 484 is private because it is associated with collector 424, which is not the root collector 420.

1 The collector implements the customer's *acquisition policy*, which determines under which
2 circumstances to add or remove resources from its collection, or change the composition of a
3 collection. The hierarchical structure of collector allows local (organization specific) policy
4 governing resources to be specified. The simplest acquisition policy (i.e., no acquisition policy)
5 is entirely request driven. That is, a resource moves between BRLS 482 (486, 488) and PMRS
6 495 through the collector hierarchy on every request and return. A more complex acquisition
7 policy might specify minimum and maximum numbers of servers at each collector. These
8 acquisition policies filter the request flow. Such an acquisition policy allows collectors to
9 maintain a variety of free pools (caches of allocated or reserved resources) for its domain, in
10 effect pre-allocating resources for predicted requests. Acquisition policy will be discussed in the
11 description of figures 7 & 8.

12 Each BRLS 482, 484, 486, 488 is associated with a collector. This association defines the scope
13 of sharing for the resource pool represented by the BRLS. For example, if an organization wishes
14 to have a private resource pool for its departments, it would have a BRLS 484 associated with
15 the collector 424 of the organization as illustrated in Figure 4. The resources would be available
16 for allocation only by customer computing environments associated with the organization (i.e.,
17 associated with a collector that has the organization's collector as an ancestor).

18 Figure 4 A shows a sub domain, Hardware 123, of the Smith Fastener company from Figure 1A.
19 For provisioning purposes Hardware has been assigned three applications, App1 440, App2 450,
20 and App3 460. In this case all three application are allowed to share the same set of resources.

21 Figure 4B shows how this decision on provisioning maps to a collector hierarchy. Col 1 470 is
22 associated with Hardware 123. In addition each application has a PMRS and a collector
23 associated with it. App1 440 has Col 1 441 and PMRS 1 442, App 2 has Col 2 451 and PMRS 2
24 452, and App3 has col 3 461 and PMRS 3 462. In this invention a PMRS is associated with at
25 most one collector.

26 Figure 5 shows an example of a computing utility with multiple root collectors 510 and 520. This
27 invention describes a BRDS 219 with a single root collector. Those skilled in the art can extend

1 this invention to multiple roots by logically partitioning the resource pools and associating a
2 BRLS for each partition with distinct root collectors. The BRLSs for the root collector could
3 manage the same resource types. Figure 5 illustrates this concept by showing a BRDS 219 with
4 two root collectors 510 and 520. Root collector 510 has BRLSs 530 and 535. Root collector 520
5 has BRLSs 540 and 545. In this case BRLS 530 could be managing the same type of resources as
6 BRLS 540. Note that in an alternative implementation, BRLSs of the same type but different
7 association can be represented as a single BRLS with constraints on where their resource
8 instances may be assigned.

9 Figure 6 illustrates an example of a collections and libraries. Each computing environment has a
10 set of resources allocated to it. These are shown only in their corresponding PMRS 660, 680,
11 and 690, not in collectors or their ancestors. Collector Col-Y 632 has an extra resource of type
12 “triangle” allocated for use by computing environment Y. The home collection of the triangle is
13 the private library BRLS-1 690. Triangles can only be allocated to PMRS X 660 and PMRS Y
14 680. Collector Col-A 630 has extra resources of type “square”, “star”, and “circle” for use by
15 computing environments Col-X 633 or Col-Y 632. They are each from a public BRLS of the
16 corresponding type 692, 694, and 696 respectively. These resource instances could be reclaimed
17 for use by collector Col-B 620. BRLS 690 in Figure 6 is private. Its resources are shared by
18 collectors 630, 632, and 633. In Figure 6 BRLSs 692, 694, and 696 are public.

19 When a resource is requested by a computing environment, the BRDS 219 goes through a three
20 stage process. The first stage is determining whether or not the acquisition policies of the
21 collector for the computing environment and its intermediate ancestor nodes, if any, would be
22 satisfied if the request is granted. The BRDS 219 may deny the request if for example the
23 resources requested would exceed that specified in the customer’s agreement with the service
24 provider. Thus the acquisition policy is a set of constraints that is taken into account in this
25 decision process. If the BRDS 219 determines that a resource can be granted, the second stage of
26 the process is to locate the instances which will satisfy the request. Available instances may be
27 found in a public or private BRLS, or in a collector. If no instances are available, then an
28 instance may be reassigned from another computing environment according to a distribution

1 policy. If the request is for a set of resources, the BRDS 219 coordinates their allocation
2 accordingly. If the request succeeds, a set of reservation tickets is returned to the PMRS.

3 The BRDS 219 will attempt to satisfy the acquisition policies of collectors, even in the absence
4 of resource requests from the PMRS, generating requests to a BRLS if necessary. It is the BRDS
5 219 that makes reservations on behalf of a collector, whether it is provoked by acquisition policy
6 or explicit resource requests.

7 Returning to the description of the collector and BRDS 219, when the start time of a reservation
8 arrives, the BRDS 219 transforms the ticket into a resource handle. If the transformation of the
9 reservation into a resource handle fails, the reservation cannot be honored. If the collector
10 holding the ticket is associated with a PMRS, the BRDS 219 delivers the handle to the collector
11 who then presents it to the PMRS. In Figure 6 if collector 632 was holding a reservation ticket
12 that became due, the BRDS would transform it into a resource handle which would be returned
13 to collector 632 who would then pass the resource handle onto PMRS 680. The PMRS 680 can
14 then configure the resource into the computing environment. Otherwise the collector holding the
15 reservation is not associated with a PMRS so the BRDS 219 delivers the resource handle to the
16 collector. In Figure 6, this would occur if a reservation were held by collector 630.

17 When a reservation expires, the BRDS 219 reclaims the resource handle associated with it, and
18 returns it to its lender (home pool or intermediate collector). The invariant maintained is that a
19 resource is checked out of a home pool or lent by an intermediate collector if and only if there is
20 a reservation for it (and presumably, a customer is being billed for it).

21 To extend a reservation, a PMRS 660 issues a request to the BRDS 219 with a selection
22 specification for the resources it is already using. The BRDS 219, if acquisition policy and other
23 commitments allow, will interact with the BRLS (692, 694, 696 or 690) to extend the
24 reservation. That is, the PMRS 60 does not communicate with the BRLS directly.

25 When a resource is returned by the PMRS 660, the acquisition policy of its associated collector

1 determines if it is returned to its parent. If the collector retains the resource, the client may
2 continue to be billed for the resource. One advantage of retaining the resource is faster
3 deployment on request. (For example, metering for the resource is already configured.) If it
4 returns the resource, the acquisition policy of its parent is similarly checked. Resources are not
5 retained at collectors associated with their home pool (BRLS); they are checked into their home
6 pool in that case. For example if a resource associated with BRLS 690 were returned to collector
7 Col-A 630, it would be returned to the BRLS 690 instead of being held at the collector.

8 Note that a collector that acts as a resource pool for a set of computing environments (i.e., an
9 intermediate node in collector hierarchy) such as 630 in Figure 6, requires the same functionality
10 as a BRLS to distribute its resources. It is essentially a library that lends borrowed resources. As
11 such, its resource collections can also be represented as BRLSs associated with it. The difference
12 between a collector and a private BRLS is that resources may be reclaimed from a collector for
13 use by its parent or other parts of the collector hierarchy (where ever the home BRLS of the
14 resource allows it). In contrast, the resources in a private BRLS can only be used by its associated
15 collector and its descendants. A collector can potentially have its own collection of every
16 resource type known to the system. Like a private BRLS, an implementation of the collector's
17 collections can be achieved within a single set of BRLSs, with constraints on where their
18 resource instances may be assigned, and additionally, the identity of the collector that holds the
19 reservation ticket for the resource from its home pool. Once a resource has been assigned to a
20 collector or a PMRS, its location within the hierarchy becomes an attribute of the resource.

21 Figure 7 illustrates an example of a process of checking acquisition policies on resource requests,
22 whether a request is generated by a PMRS or as a result of a collector's acquisition policy. The
23 process starts in box 710. We first check 720 to see if the request is initiated by a PMRS. If it is
24 730 we set the current collection to be the PMRS so that acquisition policy checking begins with
25 the collector associated with that PMRS. Otherwise 740 we set the current collection to be the
26 requesting collector's parent. This is because the request was initiated by a collector whose
27 acquisition policy requires more resources. Next we check 750 to see if we are at a root collector.
28 If the policies are satisfied all the way to the root collector, then the resources can be obtained,

1 and flow passes 780 to a procedure which attempts to locate available resources. Otherwise we
2 check 760 see if the request would violate the acquisition policy of the current collection. If the
3 policy would be violated then we check 765 to see if the requester is a PMRS. If it is the request
4 is denied 795. Otherwise, it may be possible to satisfy the request using resources already
5 allocated to the collector whose acquisition policy would be violated and its descendants, and
6 arbitration 790 is called on the part of the sub tree rooted at the collector whose acquisition
7 policy would have been violated. Returning to the check at 760, if the policy would not be
8 violated at the current collection, then we set 770 current collection to the parent of the
9 collection that was just examined and continue checking at 750.

10 Figure 8 illustrates an example of a process of locating available resource:. The process starts in
11 810. First the starting point of the search is determined. If the request originated from a PMRS
12 820, the starting collection for the search is the collector associated with the PMRS 827.
13 Otherwise, it is the parent of the collector from which the request originated 825. Since a
14 resource request may specify a set of resources of different types, a list of the descriptions of the
15 desired resources is constructed 830 , one element per resource type per instance. Those skilled
16 in the art realize that there are multiple ways to form this list, such as one element per resource
17 type with multiple elements of the same type being tracked within the element. For each element
18 in the list, three alternatives are checked for resources: a private BRLS associated with the
19 collector, the collector's collection, or a public BRLS (when the search reaches the root). These
20 checks are repeated at each collector from the starting collector to a root. This is done by
21 removing 840 the first requested resource RR, from the list of required resources. Beginning in
22 step 850 with the starting collection, we check 860 to see if there is a BRLS with an available
23 resource of type RR. If not we then check 863, to see if the collector has a resource of type RR.
24 If not we check 865, to see if we are at the root collection. If we are not at the root, then we move
25 867 to the parent of the current collection and continue checking at 860. If we are at the root,
26 then not all of the required resources are available and arbitration 870 will be used to determine
27 if the resources can be made available. If check 860 or check 863 finds the resource, then
28 information on the resource is added 880 to the reserve list. Next we check 890 to see if all
29 requested resources have been located. If they have not been located then we get the top element

off of the list 840 and continue looking for resources. If all resources have been located, then reservation requests 895 are made for the located resources and the result of the reservation request is returned as the end of the process to locate available resources.

If the reservation fails, arbitration will be called with the root collector as scope. Arbitration is the subject of a cross referenced application entitled "Apparatus for Arbitration in a Computing Utility System" having docket number YOR920030589US1, the present invention treats arbitration as a "black box". If arbitration succeeds, the resources are reserved and tickets are returned to the requester. Otherwise, the request is denied. A successful arbitration may require resources to be reclaimed from other domains. The constraints imposed by the collection managers, associated libraries, and acquisition policies are input to the arbiter.

The BRDS provides the Resource Management operations listed below.

Resource Management Operations

Figure 9 illustrates an example of the interactions between the BRDS 219 and the PMRS 920. There is one set of operations 930 for the BRDS to make requests of the PMRS and another set of operations 940 for the PMRS to make request of the BRDS. Allocation of resources involves five operations: the PMRS can request, retain, or return resources 940, and the BRDS can reclaim or offer resources 930. When a PMRS 920 needs additional resources or has excess to return it uses the RequestResource and ReturnResource interfaces 940. When the PMRS wants to keep resources beyond the duration specified on the initial request, it uses RetainResource 940. When resources that were previously reserved become available the BRDS 219 uses the DeliverResource interface 930 to give the resources to the PMRS 920. The reclaim resource interface 930 is used by the BRDS whenever it needs to remove resources from a PMRS. For example, one case would be revoking previously allocated resources so that they can be reallocated somewhere else. The OfferResource interface 930 is used to make available additional resources that were not explicitly requested by the PMRS 920 but are allowed with the policies of its computing environment.

1 **RequestResource (R, start-time, duration)**

2 This request is issued by the PMRS 920 to the BRDS 219 on behalf of a customer
3 computing environment to request additional resources. The BRDS 219 can refuse the
4 request.

5 Inputs:

6 R[], a vector of structures, one element for each resource type. Each
7 structure contains the number of instances of the type and optional
8 selection criteria. The selection criteria may include resource handles to
9 specify particular instances.

10 Start time could be immediate or for an advance reservation

11 Outputs:

12 Return code indicating success or failure

13 R'[], an optional return. On failure, may be set to indicate resources
14 available. The elements of R' are resource reservation tickets, one per
15 resource instance, which the PMRS 920 may redeem to acquire the
16 additional resources.

17 Start-time', an optional return. On failure, may be set to indicate when
18 resources are available. Applies to R' if set, R otherwise.

19 **ReturnResource (R)**

20 This request is issued by the PMRS 920 to the BRDS 219 on behalf of a customer
21 computing environment to return resources. The BRDS 219 accepts the resources.

22 Inputs:

23 R[] a list of resource identifiers (handles or reservation tickets) for the
24 resources being returned. These resources could be currently allocated
25 resources or reservations.

26 Outputs:

27 Return code indicating success or failure

1 **Retain Resource(R, end time)**

2 This request is issued by the PMRS 920 to the BRDS 219 on behalf of a customer
3 computing environment to retain resources (extend or modify the reservation of
4 resources) for a different amount of time from the previously committed time. The BRDS
5 219 may accept or reject the request.

6 Inputs:

7 R[] a list of resource identifiers (handles or reservation tickets) for the
8 resources being retained. These resources could be currently allocated
9 resources or reservations.

10 End time: is the requested new end time for the resources in the list. If the
11 end time is not in the future the request is rejected.

12 Outputs:

13 Return code indicating success or failure

14 **ReclaimResource(R)**

15 This request is issued by the BRDS 219 to a PMRS 920 to revoke the allocation
16 or reservation of a set of resources. The PMRS 920 relinquishes the resources.

17 Inputs:

18 R[], a vector of structures, one element for each resource type. Each
19 structure contains the number of instances of the type and optional
20 selection criteria. The selection criteria may include resource handles or
21 reservation tickets to specify particular instances. The computing
22 environment may select the resource instances to be reclaimed according
23 to the criteria.

24 Outputs:

25 R'[] a list of resource identifiers (resource handles or reservation tickets)
26 being returned. R' may be a superset of R which includes resources that
27 would be rendered useless because of dependencies on reclaimed
28 resources.

1 **OfferResource (R, start-time, duration)**

2 This request is issued by the BRDS 219 to a PMRS 920 to signal the availability
3 of additional resources for this computing environment. The PMRS 920 may accept or
4 decline the offered resources. To accept the resources the PMRS 920 issues a
5 RequestResource request.

6 Inputs:

7 R[] a vector containing the number of instances of each resource type
8 available

9 Start time could be immediate or for an advance reservation

10 Outputs:

11 None

12 **DeliverResource (R)**

13 This request is issued by the BRDS 219 to a PMRS 920 to indicate that previously
14 reserved resources have been allocated for the computing environment.

15 Inputs:

16 R[] a list of tuples, each consisting of a reservation-id and corresponding
17 resource-handle.

18 Outputs:

19 None

20 To acquire and distribute resources to customer computing environment, the BRDS 219 invokes
21 the operations of the BRLS 910, 912, 914, 916 described earlier.

22 Figure 10 illustrates an example of the interactions between the PMRS 1010 and the BRDS 219
23 during normal operations. The PMRS in Figure 10 represents a group of Linux servers which
24 may be of variable size. Some event causes the PMRS 1010 representing the Linux server to
25 realize that it needs additional resources, such as an xSeries Linux server and an additional IP
26 address. It requests the resources from the BRDS 219 using RequestResource 1011. The BRDS

219 uses the methods described in figures 7 and 8 to allocate and reserve the resources. The reservation tickets are returned to the PMRS 1010 in the response to RequestResource 1011. When the resources become available the PMRS 1010 will be informed using DeliverResource 1012. The detailed interactions inside of the PMRS 1010 are described in an application entitled "Componentized Automatic Provisioning and Management of Computing Environments for Computing Utilities", having docket number YOR920030588PCT1.

Figure 11 illustrates an example of a process of adding a new customer (or service) to an existing hosting infrastructure. This process is analogous to adding a new resource to an existing service. This invention assumes that there is a set of operations that allows organizations to subscribe to the services of the service provider, and further, to add computing environments in groups for resource sharing within an organization. At a minimum, to add an organization certain information from the agreement between the organization and service provider, such as the start time and duration of the agreement and acquisition or other policies is supplied. For specific computing environments, the computing environment type and settings are also supplied. During normal operations the policies controlling a computing environment can be modified by the addition of new policies, removal of existing policies, or the updating of existing policies. In addition, this invention assumes that there are operations to remove domains and organizations when their agreements terminate. To add a new domain the client specifies the representation of the domain (including sub domains), policies, and a list of computing environments to provisioned for the domain.

In Figure 11 the computing environment type is a web site. The BRDS 219 has a single root collector. There are BRLSs for network dispatcher license 1130, Linux servers (hardware and software) 1132, and WebSphere license 1134. When the new customer request 1105 arrives, the new customer (NC) 1160 component of the BRDS 219 creates a new collector 1150 for it, (indicated by dotted lines) and a place holder for the customer's information until the PMRS 1180 for the new collection is fully operational. The new customer request 1105 specifies the exact composite resource that will be instantiated. After creating the new collector the NC 1160 sends a Realize 1111 request to the PMRS 1180 of the correct type. Realize is used by the BRDS

219 to get a set of possible plans for constructing a resource on the existing infrastructure when a new computing environment is requested. Realize is defined in detail in a separate invention entitled "Componentized Automatic Provisioning and Management of Computing Environments for Computing Utilities", having docket number YOR920030588PCT1. At the end of the Realize 1111 request the BRDS 219 has list of possible sets of resources that can be used to build the computing environment of the specified type. All of these options are returned to the NC 1160 object which instantiates a collector 1150 to manage the new instance and request that the collector 1150 to negotiate for the resources. The newly instantiated collector 1150 (BRDS) analyses the options against the available resources, and policies and negotiates for resources with the libraries BRLSs 1130, 1132, 1134. It then schedules the construction of the service at the requested time assuming all of the resources can be acquired within policy. (If the resources cannot be acquired the new customer 1105 request fails.) When the start time of the resource reservations arrives, the BRDS 219 obtains the resources and delivers the resource handles to NC 1160, which sends a Build 1112 request to the PRS 1110. The Build request includes all required policies. One type of policy that is supplied is the acquisition policy, other required policies can be resource type specific and would be supplied at this time. It should be noted that a PMRS 1180 is the logical association of a PRS 1110 and an MRS 1170. The details of the PMRS are described in an application entitled "Componentized Automatic Provisioning and Management of Computing Environments for Computing Utilities", having docket number YOR920030588PCT1. As a result of the Build 1112 request the PRS 1110 creates the MRS 1170 that will be associated with the computing environment. When this process completes the PRS 1110 returns the identifier of the new MRS 1170 to the NC 1160, the NC 1160 updates the collector 1150 with the identity of the root MRS 1170 for the computing environment.

Once the collector has the handle it is now associated with the newly instantiated service which it manages, as shown by the dashed line in Figure 11. In an alternative embodiment, instead of having the NC send a build 1112 request to the PRS, the NC could send a DeliverResource request and the PMRS could realize that the first deliver resource request implies build.

In an alternate embodiment, the function of the NC can be combined into the collector so that the

BRDS could instantiate the new collector when the new customer request arrives. The new collector would, as part of its initialization, go through all of the steps assigned to the NC. At the completion of those steps the NC would be done and the service operational. It is also possible that the hosted environment might not have the ability to satisfy the request, in which case the request for a new customer would be rejected. As indicated previously, this hosted environment has the ability to acquire and utilize resources that may be in some other hosted environment. Consequently, if this hosted environment does not have all of the resources required and it can obtain the resources from some other hosted environment with an acceptable set of conditions then the request for a new customer can be accepted instead of rejected. Resources obtained from another hosted environment would still be represented as a BRLS in this environment. When a hosted environment contains BRLSs that come from multiple hosted environments, the BRLSs are typed so that acquisition policies can be supplied to manage the use of external versus internal resources.

Figure 12 shows an example of a distribution for provisioning purposes of a set of computing environments to Smith Fastener company 120. Smith Fastener had to decide which services they wanted to utilize within their organization and which part of the organization would hold (or subscribe to) each service. Smith also had to decide how to allocate resources to their domains and what the allocation and distribution policy would be for each domain. These provisioning decisions are orthogonal to which employees of Smith Fastener get to use the services that are provisioned. The guiding principle in producing the domain and computing environment tree is that a domain is only inserted where there is at least one computing environment to be provisioned, at least two sub domains to be established, or there is at least one computing environment and one sub domain. This principle prevents long strings of domains with nothing other than one sub domain. Smith Fastener Company 120 has decided that it wants to provision and control two applications at the corporate level application 1 and application 5. An accounting application, application 12 is to be used by corporate and each division Research 121, Finance 122, Hardware 123, and Marketing 124. The hardware division has decided that each product unit Bolts 127, Screws 128, and Hinges 129 should use their own instance of the same engineering design application 9. The Hinge unit 129, is also experimenting with two additional applications

1 application 10 and application 11. Marketing has decided that each of its regions Region 1, 125
2 and Region 2, 126 should use their own instance of the same application, application 3, for
3 managing marketing activities. Region 2, 126 of marketing is divided into two separate areas
4 North 130 and South 131. Each of these regions uses their own instance of application 4 to
5 collect information to feed into application 3 at the regional level. The resource acquisition and
6 distribution policies of Smith Fastener 120 assure that each of its major domains Research 121,
7 Finance 122, Hardware 123, and Marketing 124 has the resources it needs to accomplish its
8 tasks. In addition Marketing 124 and Finance 122 can have resources reallocated to them from
9 other divisions, but not each other, during their peak end of quarter and end of year times. But
10 they are not allowed to take so many resources as to disable the other divisions.

11 When all of these decisions are mapped to a domain structure with computing environments it
12 looks as follows: The root domain 120, has three applications App01 150, App12 151, and App
13 05 152. The root domain 120 also has four sub domains Research 121, Finance 122, Hardware
14 123, and Marketing 124. The sub domains of Research 121 and Finance 122 each use a single
15 application. In this case they each have separate instances of the same application App12 155 and
16 App12 156. The present invention does not prohibit the sharing of a single instance of an
17 application between domains from a user perspective. For example in Figure 12 App01 could be
18 used by both Research and Finance as well as the root domain. However, from a provisioning
19 (and management of provisioning) perspective each instance of an application has a single point
20 of provisioning within the organization. Returning to the detailed description of the figure. The
21 Hardware domain 123 has three sub domains and a single application associated with it. The sub
22 domains of Hardware 123 are Bolts Unit 127, Screw Unit 128, and Hinge Unit 129. The
23 application associated with hardware is App12 157. The Marketing domain 124 has two
24 applications and two sub domains associated with it. The application are ,App12 153 and App03
25 154 and the domains are Region 1, 125 and Region 2, 126. Each of those regions has its own
26 instance of App03. Region 1, 125 uses App03 161 and Region 2, 126 uses App03 158. Region 2,
27 126 is further divided into North 130 and South 131 and each sub domain of Region 2, 126 has
28 its own copy of App04 North uses App 04, 163 and South uses App04, 162.

Figure 13 shows the collector hierarchy that results from the computing environment and domain structure of Figure 12. To produce the collector hierarchy we start at the root domain and insert a collector. Next we examine the root domain of the domain tree to see if there is only one computing environment, if there is then we insert a PMRS for this computing environment and we are done. Otherwise, the root domain has more than one computing environment or sub domain. For each computing environment of the root domain we insert a collector and PMRS. The remaining sub domains are placed on the domain processing list. This algorithm is applied to each domain on the domain processing list until it is empty. Basically this algorithm is used breadth first search starting at the root domain of the domain and computing environment tree. When this algorithm is applied to Figure 12 the collector hierarchy shown in Figure 13 results. PMRSs are typed, in Figure 13 the last two digits of the PMRS name are its type so PMRS01 is a PMRS which manages an instance of App01 from Figure 12. Returning to a more detailed description of Figure 13, Collector SFC 1320 is inserted because Smith Fastener Company 120 in Figure 12 has more than one application. In Figure 12 Smith Fastener 120 has three applications. Therefore applications 1, 12, and 5 a collector PMRS structure was inserted. For application 1 this is collector 1352 and PMRS01 1374, for application 12 this is collector 1351 and PMRS12 1373, and for application 5 this is collector 1350 and PMRS05 1371. Next we apply the algorithm to Smith's sub domains. We notice that two of the sub domains in Figure 12 Research 121 and Finance 122 only have a single application, application 12. Consequently for each of these sub domains a collector and PMRS structure is inserted into the hierarchy. For Research 121 in Figure 12 collector 1321 and PMRS12 1370 is inserted and for Finance 122 in Figure 12 collector 1322 and PMRS12 1372 is inserted. In Figure 12 Hardware 123 and marketing 124 have multiple sub domains and application. Next we apply the algorithm to the sub domains of Hardware and Marketing. For hardware collector 1323 is inserted and for Marketing collector 1324 is inserted. In Figure 12 hardware has three sub domains Bolts 127, Screws 128 and Hinges 129 and one application, application 12 associated with it. So a PMRS collector hierarchy is added to collector 1323 for the application. In Figure 12 two of Hardware sub domains have only a single application associated with them Bolts 127 and screws 128. In Figure 13 a collector PMRS structure is associated with each of these sub domain. For Bolts it is collector 1327 and PMRS09 1376 and for screws it is collector 1328 and PMRS09 1377. In Figure 12, the third unit

1 of Hardware 123, Hinges 129 has three applications associated with it so in Figure 13 collector
2 1329 is inserted. In Figure 12 marketing has two applications and two sub domains. These
3 application are App12 153 and App03 154 and the sub domains are Region 1, 125 and Region 2,
4 126. In Figure 13 we have inserted a collector PMRS hierarchy for each of the applications.
5 These are collector 1354 and PMRS12 1378 for application 12 and collector 1355 and PMRS03
6 1379 for application 3. In Figure 12 one sub domain, North 130 has only a single application
7 associated with it, App04 163. In Figure 13 a collector PMRS hierarchy is inserted. This is
8 collector 1330 and PMRS04 1384. In Figure 12 the other sub domain of marketing, Region 2
9 126 has an application and two sub domains. Consequently in Figure 13 a collector is inserted
10 collector 1326. Returning to Hinges 129 in Figure 12, it has three applications associated with it
11 App11 164, App09 165, and App10 166. In Figure 13 A collector PMRS hierarchy is inserted for
12 each of these applications. These are collector 1356 and PMRS11 1381 for application 11;
13 collector 1357 and PMRS09 1382 for application 9; and collector 1358 and PMRS10 1383 for
14 application 10. Finally we examine Region 2 126 in Figure 12. It has one application App03 158
15 and two sub domains. In Figure 12 each of its sub domains North 130 and south 131 has an
16 instance of application 4. In Figure 13 we insert a PMRS collector hierarchy for application 3,
17 collector 1359 and PMRS03 1386. For the sub domains we also insert a collector PMRS
18 hierarchy for each one because they only have one application. This is collector 1330 and
19 PMRS04 1384 for North and collector 1331 and PMRS04 1385 for South. This completes
20 construction of the collector computing environment hierarchy which will be used to provide
21 hierarchical management of resources by this invention.

22 The collector tree is designed to allow maximum flexibility in provisioning services to a
23 customer. It allows private BRLSs to be associated with each collector in the tree. Each collector
24 can have its own acquisition and distribution policies. This allows the company, in this case
25 Smith Fastener, the ability to more finely control the distribution of its computing resources. If
26 all acquisition and distribution policies are identical and all resources are held in Public BRLSs
27 then this tree is unnecessary and a flat tree would be sufficient. By using a hierarchy, a finer level
28 of control is obtained over the distribution and use of resources within an organization.

1 The present invention is useful with the inventions described in the above cross-referenced
2 applications entitled "Componentized Automatic Provisioning and Management of Computing
3 Environments for Computing Utilities", and entitled "Apparatus for Arbitration in a Computing
4 Utility System". One example of a use is to combine the inventions to provide on-demand
5 services to a set of customers.

6 Variations described for the present invention can be realized in any combination desirable for
7 each particular application. Thus particular limitations, and/or embodiment enhancements
8 described herein, which may have particular advantages to a particular application need not be
9 used for all applications. Also, not all limitations need be implemented in methods, systems
10 and/or apparatus including one or more concepts of the present invention.

11 The present invention can be realized in hardware, software, or a combination of hardware and
12 software. A visualization tool according to the present invention can be realized in a centralized
13 fashion in one computer system, or in a distributed fashion where different elements are spread
14 across several interconnected computer systems. Any kind of computer system - or other
15 apparatus adapted for carrying out the methods and/or functions described herein - is suitable. A
16 typical combination of hardware and software could be a general purpose computer system with
17 a computer program that, when being loaded and executed, controls the computer system such
18 that it carries out the methods described herein. The present invention can also be embedded in a
19 computer program product, which comprises all the features enabling the implementation of the
20 methods described herein, and which - when loaded in a computer system - is able to carry out
21 these methods.

22 Computer program means or computer program in the present context include any expression, in
23 any language, code or notation, of a set of instructions intended to cause a system having an
24 information processing capability to perform a particular function either directly or after
25 conversion to another language, code or notation, and/or reproduction in a different material
26 form.

1 Thus the invention includes an article of manufacture which comprises a computer usable
2 medium having computer readable program code means embodied therein for causing a function
3 described above. The computer readable program code means in the article of manufacture
4 comprises computer readable program code means for causing a computer to effect the steps of a
5 method of this invention. Similarly, the present invention may be implemented as a computer
6 program product comprising a computer usable medium having computer readable program code
7 means embodied therein for causing a function described above. The computer readable program
8 code means in the computer program product comprising computer readable program code
9 means for causing a computer to effect one or more functions of this invention. Furthermore, the
10 present invention may be implemented as a program storage device readable by machine,
11 tangibly embodying a program of instructions executable by the machine to perform method
12 steps for causing one or more functions of this invention.

13 It is noted that the foregoing has outlined some of the more pertinent objects and embodiments of
14 the present invention. This invention may be used for many applications. Thus, although the
15 description is made for particular arrangements and methods, the intent and concept of the
16 invention is suitable and applicable to other arrangements and applications. It will be clear to
17 those skilled in the art that modifications to the disclosed embodiments can be effected without
18 departing from the spirit and scope of the invention. The described embodiments ought to be
19 construed to be merely illustrative of some of the more prominent features and applications of the
20 invention. Other beneficial results can be realized by applying the disclosed invention in a
21 different manner or modifying the invention in ways known to those familiar with the art.